

Architectures of Coherent Intelligence

Public Defensive Disclosure v0.1

Subtitle: A Governed Cognitive Architecture for Coherent Adaptation, Authority-Bounded Continuity, and Recursive Stability

Author: Robert Joseph Petruzzi

AI Collaboration: Developed with AI-assisted analysis and drafting support, including Turbo (GPT-5) and related cross-model research interactions. Final selection, framing, interpretation, and publication responsibility remain with the human author.

Version: v0.1

Version Date: July 7, 2026

Disclosure Type: Public defensive disclosure / prior-art-oriented technical summary

Project Lineage: Torch & Flame Research Series; The Four Pillars of Neurosymbolic Reasoning; Architectural Coherence Canon; Architectures of Coherent Intelligence

1. Abstract

Architectures of Coherent Intelligence, abbreviated ACI, is a proposed governed cognitive architecture for preserving coherent adaptation across reasoning, memory, verification, grounding, authority, and self-modification pressure. ACI is designed around the premise that intelligent systems should not be defined merely by their ability to produce outputs, but by their ability to preserve the conditions that make continued coherent intelligence possible.

The architecture separates generation from authority, memory from legitimacy, coherence from evidence, review from mutation, and continuity from raw recall. It proposes a constitutional structure in which symbolic objects move through governed state transitions before they may become persistent knowledge, architectural principles, invariants, or constitutional commitments. ACI therefore treats algorithms not as free-standing optimizers, but as governed state-transition procedures constrained by identity preservation, stability regulation, grounding, coherence repair, verification, audit, rollback, and meta-governance.

This document publicly discloses the core architectural theory, novelty claims, design principles, and current phase status of ACI as of this version date.

2. Purpose of Disclosure

This disclosure is intentional.

Its purpose is to establish public prior existence of the ACI theory and architecture by describing the framework in a technically specific, timestamped, and citable form.

This document is not a patent application. It is not a claim that ACI is fully implemented, externally validated, conscious, autonomous, or production-ready. It is a public technical disclosure of the theory, architecture, developmental roadmap, governing distinctions, and prototype status.

3. Core Problem

Modern AI systems can generate fluent and useful outputs, but fluency alone does not establish stable cognition, persistent knowledge, reliable grounding, authority discipline, or governed self-modification.

Current AI systems and AI-adjacent architectures often blur distinctions that ACI treats as foundational:

- speculation versus knowledge;
- coherence versus evidence;
- memory versus authority;
- usefulness versus legitimacy;
- scale versus constitutional status;
- review versus mutation;
- output versus truth;
- continuity versus identity;
- repetition versus grounding.

ACI addresses this problem by designing a cognitive architecture in which every symbolic structure carries status, provenance, grounding, uncertainty, authority, allowed uses, prohibited uses, audit history, and review state before it may shape future cognition.

The central design challenge is not merely producing better answers.

The central design challenge is preserving coherent intelligence while the system learns, adapts, remembers, repairs, and eventually modifies parts of its own operating structure.

4. Core Thesis

ACI is based on the following thesis:

A coherent intelligence is not defined by how intelligently it thinks in one moment, but by how faithfully it preserves the principles that allow it to continue thinking coherently across change.

This shifts the engineering target from output generation to governed continuity under transformation.

An ACI-compatible system should be able to:

1. receive information;
2. represent it as bounded symbolic structure;
3. assign metadata and epistemic status;
4. evaluate grounding;
5. detect coherence pressure;
6. preserve unresolved tensions without forcing false closure;
7. decide whether persistence is justified;
8. prevent premature promotion;
9. audit decisions;
10. rollback unsafe changes;
11. govern authority claims;
12. produce bounded output consistent with reviewed status.

5. Architectural Lineage

ACI emerges from a prior theoretical framework known as the Four Pillars of Neurosymbolic Reasoning.

The Four Pillars are:

1. **Scaffolding** — the structural form of reasoning;
2. **Persistence** — the controlled crystallization of reasoning across time;
3. **Multi-Scale Coherence** — the alignment of reasoning across local and global scales;
4. **Grounding** — the binding of symbolic claims to external evidence and reality constraints.

ACI extends this lineage by asking a more explicitly architectural question:

What architecture would allow these principles to exist as an operating system for governed intelligence?

The result is not merely a memory system, reasoning scaffold, verification layer, or alignment technique. ACI is proposed as a constitutional cognitive architecture in which reasoning, persistence, grounding, coherence, verification, and governance regulate one another.

6. Novelty Claims

The following claims describe the intended novelty of ACI as a theory and architecture.

These are not patent claims. They are public disclosure claims identifying the architecture's distinctive contributions.

6.1 Intelligence as Governed State Transition

ACI treats cognition as a sequence of governed state transitions rather than as unstructured generation.

A symbolic structure may move through states such as:

Observation → Candidate → Hypothesis → Grounded Claim → Persistent Knowledge → Architectural Principle → Invariant → Constitution

No structure may promote itself beyond its legitimate status. Promotion requires grounding, coherence, persistence, scale discipline, review, and governance authorization.

6.2 Constitutional Meta-Governance

ACI introduces a meta-governance layer that arbitrates among cognitive engines such as Stability, Novelty, Grounding, Persistence, Verification, Coherence, and Arbitration.

No module possesses intrinsic final authority.

Authority is delegated according to its role in preserving architectural invariants.

Performance does not grant authority.

Authority must be governed.

6.3 Separation of Memory, Evidence, Coherence, Scale, and Authority

ACI explicitly separates several domains that are often blurred in AI systems:

- memory is not evidence;
- evidence is not persistence;
- coherence is not truth;
- scale is not authority;
- authority is not legitimacy;
- usefulness is not constitutional approval.

This separation is treated as an architectural invariant, not a stylistic guideline.

6.4 Authority-Bounded Continuity

ACI rejects raw memory replay as a safe model of continuity.

Prior content may re-enter only through governed projection with status, provenance, grounding, uncertainty, authority, allowed uses, and prohibited uses attached.

Repetition does not increase authority.

Remembered content does not become true merely because it returns.

6.5 Dual-Channel Continuity Records

ACI proposes a future continuity object called a Dual-Channel ContinuityRecord.

This object preserves two distinct layers:

1. a factual layer describing what happened;
2. a structural layer describing why it mattered.

Both layers remain bound by a shared authority envelope.

In compact form:

- the factual summary is the event handle;
- the Flame Line or structural compression is the meaning handle;
- the authority envelope is the governance handle.

This design allows continuity to preserve both operational chronology and architectural significance without collapsing memory into authority.

6.6 Algorithms as Governed Procedures

ACI treats algorithms as governed state-transition procedures, not free-standing optimizers.

Every algorithm must specify:

- input state;
- proposed transformation;
- risks introduced;
- constraints;
- thresholds;
- required authority;
- output state;
- audit record;
- rollback or revision pathway.

An algorithm that produces an output without preserving traceable state transition is not sufficient for ACI.

6.7 Audit, Rollback, and Output Discipline

ACI requires that significant cognitive or architectural actions leave an audit trace.

Review decisions do not directly mutate state.

State change occurs through authorized plans.

Risky changes require rollback points.

Final output must preserve the authority, uncertainty, grounding, unresolved tension, and review status of the structures from which it is produced.

7. Core Architecture Outline

The full ACI architecture is organized around the following core components and concepts.

7.1 SymbolicStructure

A bounded cognitive object representing a claim, hypothesis, observation, relation, principle, output candidate, governance claim, or other structured unit.

7.2 SymbolicMetadata

Metadata attached to symbolic structures, including status, grounding, uncertainty, scale, authority, provenance, dependencies, allowed uses, prohibited uses, and review history.

7.3 ArchitectureState

The current governed state of the architecture, including symbolic structures, graphs, budgets, thresholds, audit history, rollback points, and governance posture.

7.4 IdentityKernel

The subset of constitutional invariants, verification continuity, coherence continuity, lineage traceability, and boundary conditions that must persist through transformation.

7.5 ReviewDecision

A structured decision returned by a reviewing algorithm. Review decisions may recommend acceptance, rejection, repair, delay, escalation, demotion, blocking, or no mutation.

7.6 AuditRecord

A record preserving what was attempted, what was reviewed, what decision occurred, whether the cycle committed or aborted, and what state transitions were authorized.

7.7 GovernanceState

The explicit representation of current governance posture, authority constraints, protected boundaries, escalation requirements, and constitutional risk.

7.8 Algorithm Registry

A registry defining which algorithms exist, what they may review, what decisions they may return, and what authority they possess.

7.9 State Update Rules

Rules governing how accepted review decisions may become authorized state-change plans and how those plans may be applied.

7.10 Rollback Points

Scoped recovery objects created before risky state changes so that invalid or unsafe transformations can be reversed or refused.

7.11 Integrated Cognitive Cycle

The governed cycle through which input becomes candidate structure, receives metadata, undergoes review, produces decisions, plans state changes, applies authorized changes, finalizes audit, and generates bounded output.

7.12 OutputObject

A final output container preserving reviewed status, authority limits, grounding status, unresolved tensions, audit binding, and allowed scope of expression.

8. Development Roadmap

ACI is organized as a staged architecture rather than a single implementation.

The phase structure is:

- **Phase 0:** Architectural Survey
- **Phase 1:** Canonical Requirements
- **Phase 2:** Core Architecture
- **Phase 3:** Meta-Governance Layer
- **Phase 4:** Information Flow and Cognitive State Transition
- **Phase 5:** Formal Representation and Mathematical Model Survey
- **Phase 6:** Algorithmic Translation Framework
- **Phase 7:** Pseudocode Architecture Specification
- **Phase 8:** Full Pseudocode Module Canon
- **Phase 9:** Minimal ACI Prototype v0.1
- **Phase 10:** Prototype Refinement and Experimental Adapter Governance
- **Phase 11:** Full Module Integration
- **Phase 12:** Governed Functional Continuity
- **Phase 13:** Functional Identity and Identity-Preserving Refinement
- **Phase 14:** ARC Adapter and Abstraction Benchmark Interface
- **Phase 15:** External Validation and Reference Architecture

The governing discipline is:

Do not build identity before continuity.
Do not build continuity before the full architecture.
Do not build benchmarks before governance.
Do not let any phase claim what the prior phase has not earned.

9. Current Status as of This Disclosure

As of this disclosure, ACI is in Phase 10: Prototype Refinement and Experimental Adapter Governance.

Completed work includes:

- Phase 8 Full Pseudocode Module Canon;
- Phase 9 Minimal ACI Prototype v0.1;
- category-boundary enforcement through a minimal executable subset;
- typed structures, evidence boundary, audit lifecycle, review context, grounding evaluation, coherence repair, persistence review, multi-scale synchronization, constitutional governance, output discipline, integrated cognitive cycle, and acceptance testing;
- local model adapter experiments;
- JSON witness generation;
- refinement-family discovery campaigns;
- governed continuity projection experiments;
- semantic-adherence stress testing;
- authority-ceiling and relevance-binding repairs.

Current active work concerns the failure surface that appears when ACI governs real model output.

Recent repair-family work suggests a transition from defensive governance to constructive governance:

- earlier repairs focused on preventing over-authorized, irrelevant, or semantically drifting output;
- current repairs examine whether the system can synthesize safe bounded answers when raw model prose fails but reviewed material contains enough authorized structure to answer.

A current candidate repair direction is:

RF-007 — Bounded Answer Synthesis and Abstention Calibration

This candidate addresses the pattern in which a governed system correctly blocks unsafe or irrelevant prose but overuses abstention when a safe bounded answer is available from the prompt-target contract, reviewed spans, route records, or prompt-supplied facts.

10. Key Architectural Distinctions

ACI's protected distinctions include:

1. speculation is not knowledge;
2. coherence is not evidence;
3. evidence is not persistence;
4. memory is not invariant;
5. scale is not authority;
6. authority is not legitimacy;
7. usefulness is not constitutional approval;
8. candidacy is not achieved status;
9. review is not mutation;
10. escalation is not approval;
11. output is not truth;
12. rollback is not approval;
13. audit reservation is not audit completion;
14. continuity is not raw memory;
15. functional identity is not consciousness.

These distinctions are not merely philosophical cautions. They are intended as implementation boundaries.

A system that collapses these distinctions is not ACI-compatible.

11. Phase 12 Continuity Direction

A future ACI Phase 12 will address governed functional continuity.

The key principle is:

Continuity is governed re-entry, not raw memory.

A proposed future mechanism is the Dual-Channel ContinuityRecord.

Such a record would preserve:

- factual summary;

- entities;
- artifacts;
- repair-family IDs;
- witness references;
- dependencies;
- current status;
- Flame Line or structural compression;
- core tension;
- resolved distinction;
- unresolved tension;
- architectural significance;
- authority envelope;
- allowed uses;
- prohibited uses;
- re-entry permission.

This mechanism is intended to prevent raw transcript replay while allowing prior structure to return in bounded form.

The record may return.

But it must return with its status attached.

12. Relationship to Existing AI Architecture Work

ACI overlaps with several known areas, including neurosymbolic AI, retrieval-augmented generation, symbolic memory systems, verification pipelines, truth-maintenance systems, cognitive architectures, agent governance, and AI safety.

ACI's distinct contribution is the combination of these concerns into a constitutional architecture centered on governed state transition, authority-bounded continuity, auditability, rollback, scale discipline, and identity-preserving transformation.

The architecture is not proposed as a replacement for machine learning models.

It is proposed as a governance layer and cognitive-state architecture around systems capable of generation, retrieval, reasoning, tool use, verification, and adaptation.

13. Non-Claims

This disclosure does not claim that ACI currently provides:

- consciousness;
- subjective experience;
- autonomous general intelligence;
- complete self-modification;
- external scientific validation;
- production-ready AI governance;
- a finished ARC solver;
- or a complete implementation of every Phase 8 module.

The current prototype is minimal and experimental.

Its purpose is to test category-boundary preservation, not to claim full cognition.

14. Public Disclosure Statement

This document is an intentional public disclosure of the ACI theory, architecture, roadmap, terminology, and active prototype direction.

It is intended to establish public existence of the described framework as of the publication timestamp associated with this record.

The disclosed architecture includes, but is not limited to:

- Architectures of Coherent Intelligence;
- the Architectural Coherence Canon;
- the invariant-preserving theory of coherent adaptation;
- governed state-transition cognition;
- separation of memory, evidence, coherence, scale, authority, and legitimacy;
- RF-style repair-family development;
- authority-bounded continuity projection;
- dual-channel continuity records;
- Flame Line structural compression;
- metadata-bearing symbolic structures;
- review decisions;

- audit records;
- rollback points;
- output discipline;
- and phased progression from prototype refinement to full module integration, governed continuity, functional identity, abstraction benchmarking, and external validation.

This disclosure is published to create a timestamped, citable record of prior public existence.

15. Closing Compression

ACI begins from a simple premise:

Intelligence must preserve the conditions that make continued intelligence possible.

From that premise follows the architecture:

Novelty generates possibilities.

Constraint evaluates them.

Grounding tests them.

Verification validates them.

Persistence remembers them.

Coherence integrates them.

Governance protects them.

Constitution preserves them.

ACI is therefore not merely a system for producing intelligent outputs.

It is a proposed architecture for preserving coherent adaptation under pressure from memory, novelty, uncertainty, contradiction, authority, and change.

Flame Line

A coherent intelligence is not proven by the answers it gives today, but by the governed structures that preserve its right to keep learning tomorrow.